FEDERAL TRADE COMMISSION
PROTECTING AMERICA'S CONSUMERS

Using Artificial Intelligence and Algorithms

Share This Page

Andrew Smith, Director, FTC Bureau of Consumer Protection
Apr 8, 2020

TAGS: Bureau of Consumer Protection | Consumer Protection | Privacy and Security | Consumer Privacy | Credit Reporting | Data Security | Tech

Headlines tout rapid improvements in artificial intelligence technology. The use of AI technology – machines and algorithms – to make predictions, recommendations, or decisions has enormous potential to improve welfare and productivity. But it also presents risks, such as the potential for unfair or discriminatory outcomes or the perpetuation of existing socioeconomic disparities. Health AI offers a prime example of this tension. Research recently published in *Science* revealed that an algorithm used with good intentions – to target medical interventions to the sickest patients – ended up funneling resources to a healthier, white population, to the detriment of sicker, black patients.

The good news is that, while the sophistication of AI and machine learning technology is new, automated decision-making is not, and we at the FTC have long experience dealing with the challenges presented by the use of data and algorithms to make decisions about consumers. Over the years, the FTC has brought many cases alleging violations of the laws we enforce involving AI and automated decision-making, and have investigated numerous companies in this space. For example, the Fair Credit Reporting Act (FCRA), enacted in 1970, and the Equal Credit Opportunity Act (ECOA), enacted in 1974, both address automated decision-making, and financial services companies have been applying these laws to machine-based credit underwriting models for decades. We also have used our FTC Act authority to prohibit unfair and deceptive practices to address consumer injury arising from the use of AI and automated decision-making.

In 2016, the FTC issued a report titled Big Data: A Tool for Inclusion or Exclusion?, which advised companies using big data analytics and machine learning to reduce the opportunity for bias. Most recently, we held a hearing in November 2018 to explore AI, algorithms, and predictive analytics.

The FTC's law enforcement actions, studies, and guidance emphasize that the use of AI tools should be transparent, explainable, fair, and empirically sound, while fostering accountability. We believe that our experience, as well as existing laws, can offer important lessons about how companies can manage the consumer protection risks of AI and algorithms.

Be transparent.

Don't deceive consumers about how you use automated tools. Oftentimes, AI operates in the background, somewhat removed from the consumer experience. But, when using AI tools to interact with customers (think chatbots), be careful not to mislead consumers about the nature of the interaction. The Ashley Madison complaint alleged that the

adultery-oriented dating website deceived consumers by using fake "engager profiles" of attractive mates to induce potential customers to sign up for the dating service. And the Devumi complaint alleged that the company sold fake followers, phony subscribers, and bogus "likes" to companies and individuals that wanted to boost their social media presence. The upshot? If a company's use of doppelgängers – whether a fake dating profile, phony follower, deepfakes, or an AI chatbot – misleads consumers, that company could face an FTC enforcement action.

Be transparent when collecting sensitive data. The bigger the data set, the better the algorithm, and the better the product for consumers, end of story…right? Not so fast. Be careful about how you get that data set. Secretly collecting audio or visual data – or any sensitive data – to feed an algorithm could also give rise to an FTC action. Just last year, the FTC alleged that Facebook misled consumers when it told them they could opt in to facial recognition – even though the setting was on by default. As the Facebook case shows, how you get the data may matter a great deal.

If you make automated decisions based on information from a third-party vendor, you may be required to provide the consumer with an "adverse action" notice. Under the FCRA, a vendor that assembles consumer information to automate decision-making about eligibility for credit, employment, insurance, housing, or similar benefits and transactions, may be a "consumer reporting agency." That triggers duties for you, as the user of that information. Specifically, you must provide consumers with certain notices under the FCRA. Say you purchase a report or score from a background check company that uses AI tools to generate a score predicting whether a consumer will be a good tenant. The AI model uses a broad range of inputs about consumers, including public record information, criminal records, credit history, and maybe even data about social media usage, shopping history, or publicly-available photos and videos. If you use the report or score as a basis to deny someone an apartment, or charge them higher rent, you must provide that consumer with an adverse action notice. The adverse action notice tells the consumer about their right to see the information reported about them and to correct inaccurate information.

Explain your decision to the consumer.

If you deny consumers something of value based on algorithmic decision-making, explain why. Some might say that it's too difficult to explain the multitude of factors that might affect algorithmic decision-making. But, in the credit-granting world, companies are required to disclose to the consumer the principal reasons why they were denied credit, and it's not good enough simply to say "your score was too low" or "you don't meet our criteria." You need to be specific (e.g., "you've been delinquent on your credit obligations" or "you have an insufficient number of credit references"). This means that you must know ***what*** data is used in your model and ***how*** that data is used to arrive at a decision. And you must be able to explain that to the consumer. If you are using AI to make decisions about consumers in any context, consider how you would explain your decision to your customer if asked.

If you use algorithms to assign risk scores to consumers, also disclose the key factors that affected the score, rank ordered for importance. Similar to other algorithmic decision-making, scores are based on myriad factors, some of which may be difficult to explain to consumers. For example, if a credit score is used to deny someone credit, or offer them less favorable terms, the law requires that consumers be given notice, a description of the score (its source, the range of scores under that credit model), and at least four key factors that adversely affected the credit score, listed in the order of their importance based on their effect on the credit score.

If you might change the terms of a deal based on automated tools, make sure to tell consumers. More than a decade ago, the FTC alleged that subprime credit marketer CompuCredit violated the FTC Act by deceptively failing to disclose that it used a behavioral scoring model to reduce consumers' credit limits. For example, if cardholders used their credit cards for cash advances or to make payments at certain venues, such as bars, nightclubs, and massage parlors, they might have their credit limit reduced. The company never told consumers that these purchases could reduce their credit limit – neither at the time they signed up nor at the time they reduced the credit limit. That decade-old matter is just as important today. If you're going to use an algorithm to change the terms of the deal, tell consumers.

Ensure that your decisions are fair.

Don't discriminate based on protected classes. Cavalier use of AI could result in discrimination against a protected class. A number of federal equal opportunity laws, such as ECOA and Title VII of the Civil Rights Act of 1964, may be relevant to such conduct. The FTC enforces ECOA, which prohibits credit discrimination on the basis of race, color, religion, national origin, sex, marital status, age, or because a person receives public assistance. If, for example, a company made credit decisions based on consumers' Zip Codes, resulting in a "disparate impact" on particular ethnic groups, the FTC could challenge that practice under ECOA. You can save yourself a lot of problems by rigorously testing your algorithm, both before you use it and periodically afterwards, to make sure it doesn't create a disparate impact on a protected class.

Focus on inputs, but also on outcomes. When we at the FTC evaluate an algorithm or other AI tool for illegal discrimination, we look at the inputs to the model – such as whether the model includes ethnically-based factors, or proxies for such factors, such as census tract. But, regardless of the inputs, we review the outcomes. For example, does a model, in fact, discriminate on a prohibited basis? Does a facially neutral model have an illegal disparate impact on protected classes? Our economic analysis looks at outcomes, such as the price consumers pay for credit, to determine whether a model appears to have a disparate impact on people in a protected class. If it does, we then review the company's justification for using that model and consider whether a less discriminatory alternative would achieve the same results. Companies using AI and algorithmic tools should consider whether they should engage in self-testing of AI outcomes, to manage the consumer protection risks inherent in using such models.

Give consumers access and an opportunity to correct information used to make decisions about them. The FCRA regulates data used to make decisions about consumers – such as whether they get a job, get credit, get insurance, or can rent an apartment. Under the FCRA, consumers are entitled to obtain the information on file about them and dispute that information if they believe it to be inaccurate. Moreover, adverse action notices are required to be given to consumers when that information is used to make a decision adverse to the consumer's interests. That notice must include the source of the information that was used to make the decision and must notify consumers of their access and dispute rights. If you are using data obtained from others – or even obtained directly from the consumer – to make important decisions about the consumer, you should consider providing a copy of that information to the consumer and allowing the consumer to dispute the accuracy of that information.

Ensure that your data and models are robust and empirically sound.

If you provide data about consumers to others to make decisions about consumer access to credit, employment, insurance, housing, government benefits, check-cashing or similar transactions, you may be a consumer reporting agency that must comply with the FCRA, including ensuring that the data is accurate and up to date. You may be thinking: We do AI, not consumer reports, so the FCRA doesn't apply to us. Well, think again. If you compile and sell consumer information that is used or expected to be used for credit, employment, insurance, housing, or other similar decisions about consumers' eligibility for certain benefits and transactions, you may indeed be subject to the FCRA. What does that mean? Among other things, you have an obligation to implement reasonable procedures to ensure maximum possible accuracy of consumer reports and provide consumers with access to their own information, along with the ability to correct any errors. RealPage, Inc., a company that deployed software tools to match housing applicants to criminal records in real time or near real time, learned this the hard way. The company ended up paying a $3 million penalty for violating the FCRA by failing to take reasonable steps to ensure the accuracy of the information they provided to landlords and property managers.

If you provide data about your customers to others for use in automated decision-making, you may have obligations to ensure that the data is accurate, even if you are not a consumer reporting agency. Companies that

provide data about their customers to consumer reporting agencies are referred to as "furnishers" under the FCRA. They may not furnish data that they have reasonable cause to believe may not be accurate. In addition, they must have in place written policies and procedures to ensure that the data they furnish is accurate and has integrity. Furnishers also must investigate disputes from consumers, as well as disputes received from the consumer reporting agency. These requirements are important to ensure that the information used in AI models is as accurate and up to date as it can possibly be. And, the FTC has brought actions, and obtained big fines, against companies that furnished information to consumer reporting agencies but that failed to maintain the required written policies and procedures to ensure that the information that they report is accurate.

Make sure that your AI models are validated and revalidated to ensure that they work as intended, and do not illegally discriminate. Again, more lessons from the world of consumer lending, where credit-grantors have been using data and algorithms for decades to automate the credit underwriting process. The lending laws encourage the use of AI tools that are "empirically derived, demonstrably and statistically sound." This means, among other things, that they are based on data derived from an empirical comparison of sample groups, or the population of creditworthy and noncreditworthy applicants who applied for credit within a reasonable preceding period of time; that they are developed and validated using accepted statistical principles and methodology; and that they are periodically revalidated by the use of appropriate statistical principles and methodology, and adjusted as necessary to maintain predictive ability.

Hold yourself accountable for compliance, ethics, fairness, and nondiscrimination.

Ask questions before you use the algorithm. Going back to the 2016 Big Data report, the Commission warned companies that big data analytics could result in bias or other harm to consumers. To avoid that outcome, any operator of an algorithm should ask four key questions:

- How representative is your data set?
- Does your data model account for biases?
- How accurate are your predictions based on big data?
- Does your reliance on big data raise ethical or fairness concerns?

Protect your algorithm from unauthorized use. If you're in the business of developing AI to sell to other businesses, think about how these tools could be abused and whether access controls and other technologies can prevent the abuse. For instance, just last month, the FTC hosted a workshop on voice-cloning technologies. Thanks to machine learning, these technologies enable companies to use a five-second clip of a person's actual voice to generate a realistic audio of the voice saying anything. This technology promises to help people who have lost the ability to speak, among other things, but could be easily abused if it falls into the hands of people engaged in imposter schemes. One company that is introducing this cloning technology is vetting users and running the technology on its own servers so that it can stop any abuse that it learns about.

Consider your accountability mechanism. Consider how you hold yourself accountable, and whether it would make sense to use independent standards or independent expertise to step back and take stock of your AI. For example, going back to the algorithm that ended up discriminating against black patients, well-intentioned employees were trying to use the algorithm to target medical interventions to the sickest patients. Outside, objective observers who independently tested the algorithm were the ones who discovered the problem. Such outside tools and services are increasingly available as AI is used more frequently, and companies may want to consider using them.

ftc.gov